



The efficient frontier for spot and forward purchases: an application to electricity

C-K Woo¹*, I Horowitz², B Horii¹ and RI Karimov¹

¹Energy and Environmental Economics, Inc., 353 Sacramento Street, Suite 1700 San Francisco, CA 94111, USA; and ²University of Florida, Gainesville, FL 32611-7169, USA

A local electricity distribution company (LDC) can reduce its exposure to the inherent risks of spot-price volatility and uncertain future demand via forward contracts. Management's problem is to determine the optimal forward-contract purchase. We propose a practical three-stage approach for dealing with the problem. The first stage determines an optimal purchase by solving a cost-constrained risk-minimization problem. The second stage derives the efficient frontier of tradeoffs between expected cost and cost risk from the first-stage solution, at various bounds on the expected cost. The optimal solution is found by melding the frontier with management's risk preferences. In the third stage, the model's parameters are estimated from data typically available to an LDC and used to determine its forward-contract purchase.

Journal of the Operational Research Society (2004) 55, 1130–1136. doi:10.1057/palgrave.jors.2601769

Published online 9 June 2004

Keywords: electricity; risk; decision analysis; nonlinear programming

Introduction

Privatization and the creation of a competitive spot market for electricity in the United Kingdom (UK) brought with it volatile pool prices whose effects risk-averse large consumers (eg industrial firms) and resellers (eg regional electricity suppliers) sought to mitigate through forward contracts.^{1–3} Deregulation in the United States (US) likewise resulted in volatile spot-market prices and the use of various hedging instruments, including in particular forward trades, by local distribution companies (LDCs) that are committed to provide electricity upon demand to their customers.^{4–6} And in both cases, when put to the empirical test, the implications of the forward-contracting process for electricity prices, the buyers, and the sellers were mixed and debatable.^{3,5}

The purpose of the present paper is not to join the latter debate. Rather, we direct our attention to one important and essentially neglected aspect of the forward-contracting issue: notably, the determination, from the perspective of LDC management, of the *optimal* forward-contract purchase. We broached this issue in a previous paper wherein we described a heuristic procedure for attacking the problem.⁷ The present paper extends our previous work by providing a formal analytic procedure that allows management to readily determine the forward-contract purchase that will be optimal from the standpoint of the LDC's customers, given the inherently uncertain spot prices and customer demand that encourage hedging in the first place.

The procedure, which has potential application in *any* sort of portfolio decision in which management might purchase a forward contract in order to hedge against price fluctuations in a spot market, comprises three stages. In the first stage, an optimal purchase is determined as the solution to a cost-constrained risk-minimization problem. The imposed constraint is an upper bound on the LDC's expected procurement cost. In the second stage, the efficient frontier that summarizes the Pareto-optimal tradeoffs between expected cost and risk, as proxied through the procurement-cost variance, is derived from the first-stage solution at various bounds on the expected procurement cost. The optimal solution is then determined from management's risk preferences and the 'risk and return' tradeoffs that it is willing to accept on behalf of its customers. Finally, in the third stage, the parameters of the model are estimated from data typically available to an LDC. By way of illustration, the data used at the estimation stage are based on those for a Florida-based municipal utility owned by the city's residents. The parameter estimates are then used to determine the optimal forward-contract purchase for a hypothetical LDC.

The optimal-purchase decision under extreme risk aversion

Consider an LDC with contractual obligations to provide electricity to its customers, upon demand. The LDC can satisfy that demand through either self-generation or spot-market purchases, or by reselling electricity that it has procured via a forward contract into which it has previously

*Correspondence: C-K Woo, Energy and Environmental Economics, Inc., 353 Sacramento Street, Suite 1700, San Francisco, CA 94111, USA.
E-mail: ck@ethree.com

entered. The length of the contract period depends upon the market. The New York Mercantile Exchange (NYMEX), for example, dropped its monthly contract in favour of daily contracts in 2001, and quarterly and annual contracts might also be available. Assuming all else constant, however, an LDC would ordinarily prefer a 6-month contract to, say, two 3-month contracts, because it has a lower transactions cost. For simplicity, we assume the LDC to have neither generation capacity nor existing power-purchase contracts, and that only a single contract period is at issue. Incorporating existing contract costs or fuel costs associated with generation adds computational complexity, without improving our understanding of the procurement problem.

The LDC that meets its entire electricity requirement of Q megawatt hours (MWh) for t_1 via spot-market purchases at an average price of $\$P/\text{MWh}$ will have an *ex post* procurement cost of

$$C = PQ \quad (1)$$

Relying exclusively on the spot market to acquire Q exposes the LDC to potentially very high costs that can be mitigated when management hedges its electricity purchases through fixed-price forward contracts. In particular, suppose management purchases q MWh at the forward price of $\$F/\text{MWh}$. Now the *ex post* cost is

$$C = P(Q - q) + Fq \quad (2a)$$

Equation (2a) shows that C converges to FQ as q approaches Q . While forward contracting can reduce the cost effects of unanticipated spot-price changes, it cannot eliminate the cost variations due to unanticipated changes in its customers' demand for electricity, demand that the LDC is contractually obligated to satisfy. Thus, when making its forward-contract decision, management must recognize that both P and Q are random variables, so that C is also a random variable. Suppose management assigns to P and Q expected values (variances) of μ_P (σ_P^2) and μ_Q (σ_Q^2), respectively, as well as a covariance of $\sigma_{P,Q}$.

When $Q \geq q$, the LDC will have to purchase $(Q - q)$ on the spot market at the price of P . Alternatively, when $Q < q$, the LDC will enter the spot market as a seller, rather than as a buyer, and the $P(Q - q) < 0$ term in Equation (2a) reduces the procurement cost.

Rewrite Equation (2a) as

$$C = PQ - Pq + Fq \quad (2b)$$

The expected procurement cost is⁸

$$\mu = (\mu_Q - q)\mu_P + \sigma_{P,Q} + Fq \quad (3)$$

Let ρ denote the correlation between P and Q . Casual observation leads us to hypothesize $\rho > 0$. Specifically, the two most common causes of *sharp* spikes in the spot price are: (1) relatively small supply reductions, say due to forced plant outages, along price-insensitive spot-market demand

curves; and (2) demand surges, say due to rising temperatures in the summer or falling temperatures in the winter, along close to full-capacity and virtually inelastic supply curves.^{9,10} We conjecture that the latter tend to be more responsible for short-term price changes than are the former. As will subsequently be seen, this conjecture, which implies $\rho > 0$, is supported by the sample data that underlie our empirical analysis. The conjecture presumes that generators are not 'gaming the market' by withdrawing capacity at critical times of near-capacity demand in order to drive up prices. California and the United Kingdom, however, may provide recent examples that call into question the universality of the presumption.

Since $\rho > 0$, $\sigma_{P,Q} = \rho\sigma_P\sigma_Q > 0$. Denote $\mu_{CE} = (\mu_Q - q)\mu_P + Fq$ as the certainty-equivalent procurement cost, or the procurement cost evaluated at the expected price and quantity. It is immediately seen that $\mu \geq \mu_{CE}$, with the strict equality holding when P and Q are uncorrelated. Moreover, $d\mu/dq = -\mu_P + F$ and $d^2\mu/dq^2 = 0$. Hence, so long as $0 < F < \mu_P$, $d\mu/dq < 0$; or, the expected procurement cost is a decreasing linear function of forward-contract purchases. When $F = \mu_P$, $d\mu/dq = 0$ and the expected procurement cost is invariant with respect to forward-contract purchases. Otherwise, the expected procurement cost is an increasing linear function of the forward-contract purchase. Why, then, would LDC management enter into a forward contract that commits it to purchases at a price of $F > \mu_P$? The answer is to reduce risk, as measured through the procurement-cost variance⁸

$$\sigma^2 = \sigma_{PQ}^2 + q^2\sigma_P^2 - 2q\sigma_{PQ,P} \quad (4)$$

Here, σ_{PQ}^2 is the variance of the product PQ , σ_P^2 is the spot-price variance, and $\sigma_{PQ,P}$ is the covariance between PQ and P . Both σ_{PQ}^2 and $\sigma_{PQ,P}$ can be expressed in rather daunting equations, containing the expected values and variances of P and Q , as well as expectations of higher-order moments.⁸ Happily, these expressions are irrelevant to the present analysis or to the illustration that follows.

Equation (4) makes it apparent that a high spot-price variance directly translates into a high procurement-cost variance. As $d\sigma^2/dq = 2q\sigma_P^2 - 2\sigma_{PQ,P}$ and $d^2\sigma^2/dq^2 = 2\sigma_P^2 > 0$, the procurement-cost variance is a strictly convex function of q that takes on its minimum value at $q = \sigma_{PQ,P}/\sigma_P^2$. Thus, the procurement-cost variance decreases with forward-contract purchases when $q < \sigma_{PQ,P}/\sigma_P^2$ and increases with forward-contract purchases when $q > \sigma_{PQ,P}/\sigma_P^2$. More critically, under extreme risk aversion, when management's sole concern is with minimizing the cost variance, it will enter into a forward contract for $q = \sigma_{PQ,P}/\sigma_P^2$, regardless of the cost! But $q > 0$ requires $\sigma_{PQ,P} = E((PQ - \mu_{PQ})(P - \mu_P)) = E(P^2Q) - \mu_P E(PQ) > 0$. The latter expectation is necessarily positive, as both P and Q are positive. Hence a minimum-variance solution with $q > 0$ requires $E(P^2Q)/E(PQ) > \mu_P$. Otherwise, there is no feasible solution.

Stage 1: the optimal purchase decision under an expected-cost constraint

Extreme risk aversion is rare. Rather, suppose that a risk-averse management sets an upper bound M on the expected procurement cost that it can tolerate, and wants to determine its forward-contract purchases so as to minimize the cost variance in light of M . We assume that the risk-averse management does not engage in speculative short selling of forward contracts. Thus, management faces the following problem:

$$\text{Minimize}_q \quad \sigma^2 = \sigma_{PQ}^2 + q^2 \sigma_P^2 - 2q\sigma_{PQ,P} \quad (5a)$$

$$\text{Subject to } \mu = (\mu_Q - q)\mu_P + \sigma_{P,Q} + Fq \leq M \quad (5b)$$

$$q \geq 0 \quad (5c)$$

Let λ denote a Lagrange multiplier and let * denote an optimal solution. Write the Lagrangian as

$$L = \sigma_{PQ}^2 + q^2 \sigma_P^2 - 2q\sigma_{PQ,P} + \lambda((\mu_Q - q)\mu_P + \sigma_{P,Q} + Fq - M) \quad (6)$$

Since σ^2 is strictly convex in q and the cost constraint is linear in q , the second-order sufficiency conditions for optimality are satisfied. The necessary first-order Karush–Kuhn–Tucker conditions may then be written as follows:¹¹

$$\partial L / \partial q = 2q^* \sigma_P^2 - 2\sigma_{PQ,P} + \lambda^*(F - \mu_P) \geq 0 \quad (6a)$$

$$\{\partial L / \partial q\} q^* = \{2q^* \sigma_P^2 - 2\sigma_{PQ,P} + \lambda^*(F - \mu_P)\} q^* = 0 \quad (6b)$$

$$\partial L / \partial \lambda = (\mu_Q - q^*)\mu_P + \sigma_{P,Q} + Fq^* - M \leq 0 \quad (6c)$$

$$\{\partial L / \partial \lambda\} \lambda^* = \{(\mu_Q - q^*)\mu_P + \sigma_{P,Q} + Fq^* - M\} \lambda^* = 0 \quad (6d)$$

$$q^* \geq 0 \quad (6e)$$

$$\lambda^* \geq 0 \quad (6f)$$

The zero forward-contract purchase decision

First consider the case where the optimal decision is not to engage in forward contracting, or $q^* = 0$. Suppose $\lambda^* = 0$, too. Substituting into Equation (6a), the necessary condition becomes $-\sigma_{PQ,P} \geq 0$, which as we have shown cannot be so. Hence, zero forward-contract purchases are incompatible with a non-binding expected procurement-cost constraint. This leads to:

Proposition 1 *As long as the cost constraint is not binding, management will move towards the minimum-variance solution of $q^* = \sigma_{PQ,P} / \sigma_P^2$.*

When, however, $\lambda^* > 0$, the bracketed term $\{\}$ in Equation (6d) must equal zero. Substituting $q^* = 0$ into that term evolves into the condition that $\mu_Q \mu_P + \sigma_{P,Q} = M$. This leads to:

Proposition 2 *Zero forward-contract purchases are compatible with a binding expected procurement-cost constraint if and only if the upper bound set on the expected procurement cost equals $(\mu_Q \mu_P + \sigma_{P,Q})$.*

Positive forward-contract purchase decisions

Consider the case where $q^* > 0$. For openers, once again suppose that $\lambda^* = 0$. Now it is the bracketed term in Equation (6b) that must equal zero. Substituting $\lambda^* = 0$ into that term and solving, we return to minimum-variance solution of

$$q^* = \sigma_{PQ,P} / \sigma_P^2 \quad (7a)$$

As stated in Proposition 1, when a positive optimal purchase decision results in a non-binding expected procurement-cost constraint, the optimal forward purchase does not depend on the forward price. This makes sense, because if ‘money is not a problem,’ then management should not worry about the forward price in deciding on the amount of forward purchase to minimize the procurement-cost variance.

Alternatively, suppose that $\lambda^* > 0$ so that the expected procurement-cost constraint is binding. In this case, setting the bracketed term in Equation (6d) equal to zero and solving results in

$$q^* = (M - \mu_Q \mu_P - \sigma_{P,Q}) / (F - \mu_P) \quad (7b)$$

Proposition 3 *A binding cost constraint and a positive forward-contract purchase require either $F < \mu_P$ and $M < \mu_Q \mu_P + \sigma_{P,Q}$, or $F > \mu_P$ and $M > \mu_Q \mu_P + \sigma_{P,Q}$. In either case, the optimal forward purchase $q^* > 0$ behaves like an input demand: conditional on μ_Q , μ_P and $\sigma_{P,Q}$, q^* would increase with the upper bound M and decrease with the forward price F .*

Returning to the bracketed term in Equation (6b) and solving for λ^* results in

$$\lambda^* = 2(q^* \sigma_P^2 - \sigma_{PQ,P}) / (\mu_P - F) \quad (7c)$$

Hence, if $\lambda^* > 0$ and $F > \mu_P$, $q^* < \sigma_{PQ,P} / \sigma_P^2$. The latter inequality is reversed when $\lambda^* > 0$ and $F < \mu_P$. This leads to:

Proposition 4 *When the forward-contract price is above (below) the expected spot price and management's expected procurement-cost constraint is binding, the LDC's forward-contract purchase is below (above) the amount that minimizes the cost variance.*

Intuitively, if forward contracting at a high price is less of a bargain relative to the spot market, management is willing

to accept the additional risk that accompanies a greater reliance on the spot market. Should forward contracting become more of a bargain relative to the spot market, management is willing to accept the additional risk that accompanies too large of a reliance on forward contracting.

Stage 2: the efficient frontier

Developing the efficient frontier

Now, F and μ_P are known values over which management has no control, except insofar as how management may attain a forward-contract price quote. As detailed in our earlier paper, one particularly appealing way of arriving at the lowest contract price is through an internet-based multi-round auction.⁷ Subject to that qualification, the only parameter that management *does* control is M , which dictates q^* and the implied μ^* and σ^{*2} .

Changes in M by management alter μ^* and σ^{*2} . The various (μ^*, σ^{*2}) combinations trace out the *efficient frontier* of Pareto-optimal solutions, and they provide the tradeoffs available to management between the expected procurement cost and the cost variance. When $F = \mu_P$, however, the expected procurement cost is fixed at $\mu = \mu_Q \mu_P + \sigma_{P,Q}$ and is also out of management's control. In that event there are no available tradeoffs and management, which has an obligation to serve, will elect the unique minimum-variance purchase decision. Otherwise, with the expected procurement cost shown on the horizontal axis and the cost variance on the vertical axis, the efficient frontier will be a strictly convex decreasing curve within the relevant range.

As regards the relevant range, when $F > \mu_P$, the lower bound on μ^* is dictated by the $M > \mu_Q \mu_P + \sigma_{P,Q}$ condition. The upper bound on μ^* is located where the expected procurement-cost constraint is no longer binding and the expected procurement cost is at the minimum-variance asymptote of

$$\mu_{MV}^* = \mu_Q \mu_P + \sigma_{P,Q} + (F - \mu_P)(\sigma_{P,Q,P}/\sigma_P^2) \quad (8)$$

When $F < \mu_P$, the upper bound on μ^* is dictated by the $M < \mu_Q \mu_P + \sigma_{P,Q}$ condition. In principle, there is no lower bound on μ^* . Management might see an opportunity to take advantage of the low-forward price in period t_0 to become a seller in the spot market in period t_1 , and not simply be forced to make spot sales because of low customer demand. As such, the LDC might turn an expected profit in the spot market, which could translate into a negative expected procurement cost. In practice, however, when $F < \mu_P$, the difference is small and transitory, and insufficient to encourage this sort of speculation.

As regards the curvature of the efficient frontier, since $\partial L/\partial M = -\lambda \leq 0$, increases in M , which increase μ^* in the relevant range, decrease σ^{*2} in that range. Thus the efficient

frontier is a decreasing curve. Further, since q^* is a linear function of M and μ is a linear function of q , μ^* is a linear function of M . But σ^2 is a strictly convex function of q , so that it is also and necessarily a strictly convex and decreasing function of M and hence of μ^* .

Selecting the optimal point on the efficient frontier

Management's problem is to select a specific (μ^*, σ^{*2}) pair from the combinations on the efficient frontier. When doing so, management is acting on the behalf of its customers for whom the LDC serves as a purchasing agent. One way or another, then, management must assess the tradeoffs between the expected procurement cost and the cost variance that are compatible with its customers' risk preferences. That assessment may initially be formalized through a risk-preference function, $V = v(C)$, and then translated into $E[V] = f(\mu, \sigma^2)$. Under risk aversion, the latter is a concave function for determining the optimal q^* at the (μ^*, σ^{*2}) combination that falls on the lowest indifference curve tangent to the efficient frontier.

In the fortuitous case where C is normally distributed and management quantifies its risk preferences via, say, $V = 1 - e^{-\lambda C}$, $E(V) = 1 - \exp(\gamma\mu + \gamma^2\sigma^2/2)$.⁸ The indifference curve for which $E(V) = 1 - e^{-Y}$ has $Y = \gamma\mu + \gamma^2\sigma^2/2$; or $\sigma^2 = 2Y/\gamma^2 - 2\mu/\gamma$. Thus, the indifference surface is a series of parallel lines with slope of $-2/\gamma$. The slope reflects the tradeoff that management is willing to accept between the expected procurement cost and the cost variance, where $\gamma > 0$ is the Pratt–Arrow measure of risk aversion. The latter will be well below unity even for highly risk-averse decision makers.¹²

Efficient frontiers have been used in this fashion, and for related purposes, in a wide variety of contexts,^{13–17} the most closely related of which are currency hedging¹⁸ and hedging foreign investment in US real estate through forward contracting.¹⁹ Using the efficient frontier for our specific purpose of determining the optimal tradeoff between cost expectation and cost variance is a variation on Markowitz's classic theme for portfolio selection.²⁰

Stage 3: parameter estimation

The parameters of the model, the means and variances and the like, may be estimated in any number of different ways, including direct assignment by management as in the Bayesian tradition, or as we suggest here via judicious use of regression analysis. To illustrate that procedure for a hypothetical LDC, we use the data collected for a small municipal utility (MU) in Florida to which we have guaranteed anonymity. During the data-collection period, management's planning period was t_1 = October 2002.

The estimation procedure

We follow a regression-based procedure that is detailed and justified elsewhere^{7,21} to obtain estimates of μ_P and σ_P^2 , denoted u_P and s_P^2 . In essence, the procedure entails estimating a spot-price regression that relates the LDC's monthly average spot-purchase prices to monthly average spot prices in May 2000–July 2002 at the major trading hubs of Entergy (Louisiana) and ERCOT (Texas), which are geographically close to Florida. For our hypothetical LDC, the spot-price regression's dependent variable is the monthly average price for the MU's historic purchases and the explanatory variables (besides the intercept) are the monthly average of daily spot prices at Entergy and ERCOT, where electricity forwards are traded. The sample period is May 2000 to July 2002.

Table 1 reports summary statistics of the MU, Entergy and ERCOT monthly average prices, the MU's monthly spot MWh purchases and costs, as well as the Augmented Dickey-Fuller (ADF) statistics to test the null hypothesis that a data series is a random walk.²² The ADF statistics indicate that all three of the price series follow random walks, suggesting the possibility of a 'spurious regression' wherein the MU price series and the Entergy and ERCOT price series may diverge over time without limit. The test for this is a cointegration test for stationary residuals. The test statistic is an ADF statistic whose critical value at the 5% significance level is equal to -3.34 .

Table 2 reports the spot-price regression results and the corresponding ADF statistic. The adjusted R^2 indicates that the estimated regression explains 84% of the MU price variance. The coefficient estimates for the Entergy and ERCOT prices are significant at the 5% level. The mean squared error is large (\$127/MWh) due to the relatively small sample size. The ADF statistic of -5.38 indicates that the estimated regression is not prone to spurious interpretation.

We use the coefficient estimates in Table 2 and the forward prices of \$24.80/MWh and \$27.80/MWh quoted on September 9, 2002 for October delivery at Entergy and

ERCOT, respectively, to obtain the October expected price of $u_P = \$43/\text{MWh}$ and variance of $s_P^2 = \$144/\text{MWh}$.⁷

Next, we apply an autoregressive method (PROC FORECAST in SAS) to estimate μ_Q and σ_Q^2 , which we denote $u_Q = 16\,275 \text{ MWh}$ and $s_Q^2 = 16\,297\,369 \text{ MWh}$, based on the MU's monthly net MWh purchases in May 2000–July 2002.⁷ Then we compute $r = 0.42$, the estimate of ρ , the correlation between the MU's monthly spot-purchase price and its monthly net purchases using the data for May 2000–July 2002. The positive estimated correlation between spot-purchase price and the MU's monthly net purchases supports our earlier conjecture. We compute $r_{PQ,P} = 0.93$, the estimate of $\rho_{PQ,P}$, the correlation between the MU's monthly spot-purchase cost and monthly spot price, using the data for May 2000–July 2002. Finally, the estimated mean of PQ is $u_{PQ} = \$722\,729$ whose estimated variance is $s_{PQ}^2 = \$225\,684 \text{ million}$.⁷

The LDC's efficient frontier and the optimal forward purchase

The optimal forward purchase q^* is based on Equations (7a) and (7b). The computation of the estimated expected procurement cost, u^* , and the estimated variance in the procurement cost, s^{*2} , is based on Equations (3) and (4).

Table 3 presents q^* , u^* , and s^{*2} under alternative assumptions on F and M . Consider the first three columns

Table 2 Spot price regression results

Independent variable	Coefficient
Intercept	4.10
Entergy price	0.80 ^a
ERCOT price	0.70 ^a
Adjusted R^2	0.84
Mean squared error	127
ADF statistic for testing H_0 : the price series drift apart without limit	-5.4^a

^aSignificant at the 5% level.

Table 1 Summary statistics for monthly average prices (\$/MWh), MU's spot MWh purchase, and MU's purchase cost (\$)

Statistics	MU price	Entergy price	ERCOT price	MU spot MWh purchase	MU spot purchase cost
Sample size	25	27	27	27	25
Mean	64.13	40.63	43.25	16 351	1 113 028
Minimum	25.09	18.90	18.35	10 289	276 207
First quartile	42.57	26.01	25.68	12 901	718 344
Median	61.76	40.65	45.13	15 180	870 171
Third quartile	88.25	50.88	50.65	20 129	1 483 377
Maximum	121.59	79.76	89.44	23 207	2 620 872
Standard deviation	27.79	17.01	20.40	3 967	626 694
ADF statistic for testing H_0 : the data series is a random walk	-1.43	-1.76	-1.17	-2.80	-164

Note: The MU price series only has 25 observations due to missing values.

Table 3 Efficient forward purchase q^* (MWh), expected procurement cost u^* (\$thousands), and cost variance s^{*2} (\$billions)

$F = \$39/\text{MWh}$			$F = \$50/\text{MWh}$		
q^*	u^*	s^{*2}	q^*	u^*	s^{*2}
76 969	400 000	236	4006	750 000	186
65 044	450 000	129	11 352	800 000	126
53 119	500 000	60	18 697	850 000	80
41 195	550 000	31	26 042	900 000	49
37 995	563 418	29	37 995	981 359	29

Note: When the cost variance is above its minimum of \$29 billion in the bottom row, the cost constraint is binding and $u^* = M$. At the minimum cost variance, u^* is computed according to Equation (8) in the text.




Figure 1 Efficient frontiers, indifference curves, and optimal pairs of expected procurement cost and cost variance. Note: Each frontier is intentionally drawn to include a small upward-sloping segment, which is beyond the frontier's relevant range, so as to indicate the minimum cost variance of \$29 billion reported in the bottom row of Table 3.

that contain (q^*, u^*, s^{*2}) for $F = \$39/\text{MWh}$, which is the forward price paid by the MU and less than $u_P = \$43/\text{MWh}$.⁷ The results confirm Proposition 5 that when $F < u_P$ and management's expected procurement-cost constraint is binding, q^* exceeds the amount that minimizes the cost variance. The pairs of (u^*, s^{*2}) fall on the efficient frontier for $F = \$39/\text{MWh}$, which is sketched in Figure 1. This frontier is intentionally drawn to include a small upward-sloping segment, which is beyond the frontier's relevant range, so as to indicate the minimum cost variance of \$29 billion reported in the bottom row of Table 3.

The fourth column of Table 3 shows the optimal purchases from which the efficient frontier for $F = \$50/\text{MWh}$ is derived. Since the forward-contracting costs have increased, this frontier lies above the one for $F = \$39/\text{MWh}$, as shown in Figure 1.

The two straight lines in Figure 1 reflect alternative risk preferences for the hypothetical LDC management. For

expository purposes only, C is assumed to be normally distributed, and management is assumed to have assessed an exponential risk-preference function with constant absolute risk aversion. Two levels of risk aversion are considered in Figure 1: $\gamma = 1.33 \times 10^{-6}$ and 2.0×10^{-5} . Because u^* is in thousands and s^{*2} is in billions of US dollars, the slope of the indifference curve for $\gamma = 1.33 \times 10^{-6}$ is -1.5 and that for $\gamma = 2.0 \times 10^{-5}$ is -0.1 . The hypothetical LDC's optimal forward-contract purchases for these two γ levels are 61 100 and 39 535 MWh, respectively.

Conclusions

The UK and the US are merely exemplars of a larger set of nations in which LDCs can satisfy some or all of their customers' electricity demands by dipping into spot markets. The downside of doing so is that real-time changes in those demands and shifts in electricity supply can result in wide swings in the spot price, swings that may have potentially disastrous financial consequences for an LDC. The April 2001 bankruptcy of Pacific Gas and Electric (PG&E), one of the largest utilities in the US, is a dramatic case in point. Had PG&E bought sufficient forward electricity as to lock in its cost of resale, it would not have incurred the large loss due to the difference between the spot purchase price and the capped resale rate. The forward price in the summer 2000 for delivery over the subsequent twelve months was around \$60/MWh, which was almost identical to the capped rate, excluding the charges for transmission and distribution services, and substantially below the average spot price of over \$250/MWh during the first five months of 2001.²³

One intriguing option available to management for lessening an LDC's reliance on the spot market with its inherent risks is to enter into a fixed-price forward contract for future electricity delivery to meet end-use customers' demands. That option raises three basic issues: (1) How to obtain the best contract price; (2) When to enter into the contract; and (3) The amount to be purchased in the contract. We describe an internet-based competitive auction procedure for dealing with the first issue in an earlier paper.⁷ The second issue remains a back-burner challenge that we hope to move to the front burner at some future date. In this paper, we have focused our attention exclusively on the third issue and described a readily implemented three-stage approach for resolving it.

At the heart of the approach is a philosophy that an LDC's procurement policy and strategy can and should be based on the solution to an optimization problem. One aspect of the problem requires management to identify the optimal set of forward purchases available to it. No single purchase quantity can be identified as optimal, because the problem is two-dimensional, with the objective of low expected procurement cost conflicting with that

of low cost risk. We have shown how to identify the optimal set by solving a series of expected-cost-constrained risk-minimization problems and summarizing the solutions in an efficient frontier that displays the Pareto-optimal tradeoffs between the expected procurement cost and the cost variance. The solutions, however, depend upon the values assigned to a set of parameters, and we have also described, albeit briefly here, the use of regression analysis to assign those values.

A second aspect of the problem requires management to quantify the tradeoffs between expected cost and cost risk that it is willing to accept on behalf of its customers. That, assuredly, is a somewhat more difficult problem, but it is one that all managers face, in one form or another, and represents one of the reasons why top management 'earns the big bucks.' Once management has bitten the bullet and quantified those tradeoffs, the optimal forward-purchase quantity is immediately determined.

References

- 1 Helm D and Powell A (1992). Pool prices, contracts and regulation in the British electricity supply industry. *Fiscal Stud* **13**: 89–105.
- 2 Powell A (1993). Trading forward in an imperfect market. *J R Stat Assoc* **103**: 444–453.
- 3 Lowrey C (1997). The pool and forward contracts in the UK electricity supply industry. *Energy Pol* **25**: 413–423.
- 4 Gersten A (1999). Hedging your megawatts. *J Account* **188**: 47–54.
- 5 Stavros R (2000). Risk management: where utilities still fear to tread. *Pub Util Fort* **138**: 40–47.
- 6 Faruqui A, Chao HP, Niemeyer JP and Stahlkopf K (2001). Analyzing California's power crisis. *Energy J* **22**: 29–52.
- 7 Woo CK, Karimov RI and Horowitz I (2004). Managing electricity procurement cost and risk by a local distribution company. *Energy Pol*, in press.
- 8 Mood AM, Graybill FA and Boes DC (1974). *Introduction to the Theory of Statistics*. McGraw-Hill: New York.
- 9 Borenstein S (2002). The trouble with electricity markets and California's electricity restructuring disaster. *J Econ Persp* **16**: 169–189.
- 10 Wolfram CD (1999). Measuring duopoly power in the British electricity spot market. *Am Econ Rev* **89**: 805–826.
- 11 Hillier FS and Lieberman GJ (1986). *Introduction to Operations Research*, 4th edn. Holden-Day: Oakland.
- 12 Horowitz I and Thompson P (1995). The sophisticated rdecision maker: all work and no pay? *Omega Int J Mngt Sci* **23**: 1–11.
- 13 Horsky D and Nelson P (1996). Evaluation of salesforce size and productivity through efficient frontier benchmarking. *Mkt Sci* **15**: 301–320.
- 14 Li SX and Huang Z (1996). Determination of the portfolio selection for a property-liability insurance company. *J Opl Res Soc* **88**: 257–268.
- 15 Ballesterio E (1998). Approximating the optimum portfolio for an investor with particular preferences. *J Opl Res Soc* **49**: 998–1000.
- 16 Kim SC, Horowitz I, Young KK and Buckley TA (2000). Flexible bed allocation and performance in the intensive care unit. *J Opl Mngt* **18**: 427–443.
- 17 Oliver RM and Wells E (2001). Efficient frontier cutoff policies in credit portfolios. *J Opl Res Soc* **52**: 1025–1033.
- 18 Eaker MR and Grant DM (1990). Currency hedging strategies for internationally diversified portfolios. *J Port Mngt* **8**: 30–32.
- 19 Ziobrowski B and Ziobrowski A (1995). Using forward contracts to hedge foreign investment in US real estate. *J Prop Eval and Inv* **13**: 22–43.
- 20 Markowitz HM (1959). *Portfolio Selection*. John Wiley & Sons: New York.
- 21 Woo CK, Horowitz I and Hoang K (2001). Cross hedging and value at risk: wholesale electricity forward contracts. *Adv Inv Anal Port Mngt* **8**: 283–301.
- 22 Davidson R and McKinnon JG (1993). *Estimation and Inference in Econometrics*. Oxford: New York.
- 23 Woo CK (2001). What went wrong in California's electricity market? *Energy The Int J* **26**: 747–758.

Received January 2003;
accepted September 2003 after one revision